CDF Computing Strategy and Progress

Ashutosh Kotwal
for the CDF Offline Group

Run 2 Computing Review
September 14, 2005

Outline

- Mission
- Software and Algorithms
- Infrastructure Development
- Operations
- Future Plans
- Summary

Mission Summary

- Rapid turnaround from taking collider data to physics results and publications
- Use distributed, shared GRID computing resources to meet growing needs for analysis
  - Use standard, supported tools for computing and data-handling
  - Minimize disruption in user interfaces to maximize physics productivity
- Finalize reconstruction and simulation software to allow collaboration to focus on data analysis
- Streamline operations to reduce personnel needs

CDF Physics Goals

- "1 fb^{-1} challenge for 2005-2006"
  - Present results using 1/fb of analysed data at winter conferences, and bulk of results using (1+) fb^{-1} at summer 06 conferences
- Follow with "2 fb^{-1} challenge" for 2006-2007

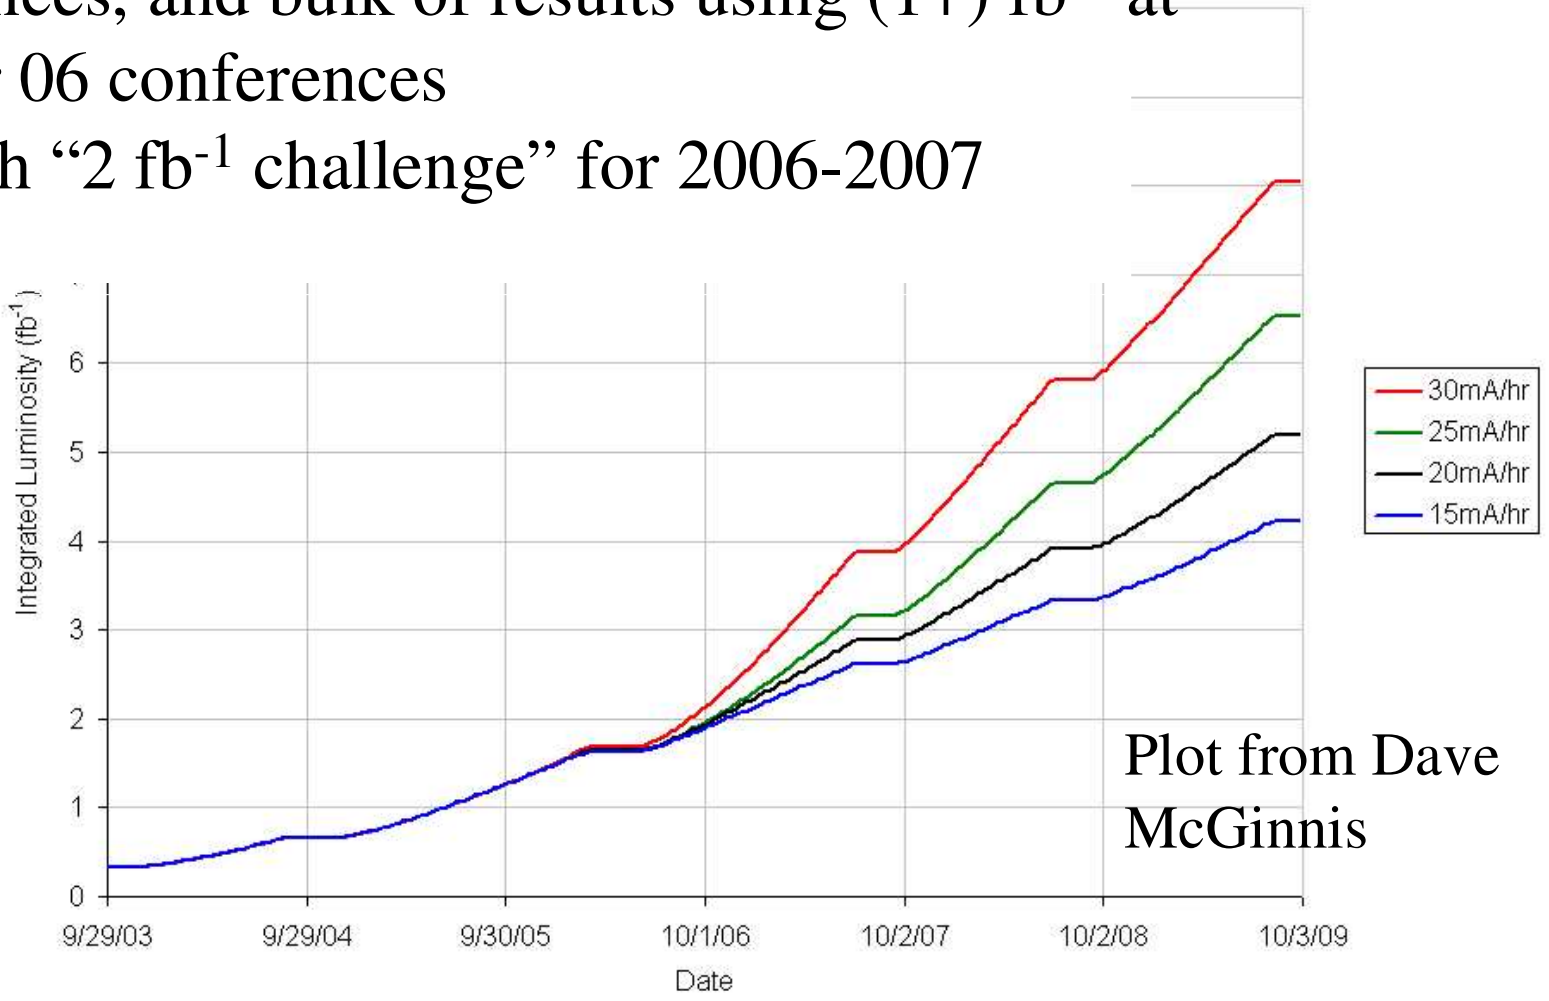

Plot from Dave McGinnis

Software and Algorithms

Software Status

- Produced new software release for 2005
- Summary of reconstruction algorithm improvements:
  - ZvertexFinder: improved efficiency and reduction in fake rate of vertex-finding
  - Main drift chamber (COT) reconstruction: recovery of missed hits
  - Fitting timing (t_0) for COT tracks and vertices
  - Improved COT dE/dx corrections
  - Time-of-Flight reconstruction improvements
  - Latest silicon detector alignment
- Summary of simulation improvements:
  - Tuned COT simulation (t_0, resolution, hit width, efficiency)
  - Plug calorimeter shower-max simulation improvements

Completed Software Projects

- Summary of infrastructure code improvements:
  - Integration of SAM into data-handling interface
  - Integration of FronTier (distributing databases)
- Switch to GCC compiler (from KAI)
  - Reconstruction and simulation code running in 'maxopt' mode in gcc v3.4.3
- Reconstruction farm and many offsite CPU farms fully operational under Scientific Linux 3.0.4
  - Migration of remaining analysis farms underway
- Summary of new code for new Run 2b detectors:
  - Timing of Electromagnetic calorimeter signals
  - Central preshower upgrade
  - New DSP code for TDC's to reduce online deadtime

Software Status

- Reconstruction code essentially stable
  - No major changes anticipated for remaining years
- Simulation code reaching high level of maturity
  - Good agreement with data
  - Simulation machinery incorporates run-dependent
    - detector configurations
    - Multiple interactions as a function of instantaneous luminosity
  - Allows straight-forward extension of Monte Carlo datasets to match collider data
- High level of operational stability (crash rate ~ 1/100M) achieved
- Conclusion: software is close to being in stable, maintenance mode

Computing Infrastructure Status

Computing Infrastructure Development

- Goal: create a uniform, GRID-enabled computing platform for all computing
  - Data processing
  - Monte Carlo production
  - User analysis
- Accomplishments in 2004-2005
  - SAM deployment
  - New reconstruction farm “SAM-Farm” deployment
  - FronTier deployment
  - Access to Fermi-Grid
  - Glide-CAF installations on LHC computing pools

Current CPU Infrastructure

- All batch CPU organized as “CAFs”, built from commodity PC's
- Use Condor batch system
- Two large CAFs at FNAL, with good connectivity to data storage, used mainly for user data analysis
- Third, smaller CAF, built this year, used for primary data reconstruction
- ~10 CAF-like installations (dCAFs) distributed around the world at collaborating institutions (in Asia, Canada, Europe, USA), providing ~45% of total CPU
  - Installing a dCAF is a documented, mature procedure
  - Sys-admin provided by local group
- SAM used for data-handling on all dCAFs
- A major goal for 2004-2005: deploy SAM on FNAL CAFs

SAM Deployment at FNAL

- Raw data logged via SAM
- All reconstructed data from 2004-2005 accessible only via SAM (~10 TB/day, ~40% of total data-handling, rest is 2004 data)
- All prior data also accessible via SAM
- User community adapting quite well to SAM
- Created CDF Help-Desk for user support
  - Manned by 5 CDF power-users, one from each physics group, and an expert
  - User questions first fielded by Help-Desk, then passed on to CDF data-handling experts and finally to CD SAM team if necessary
  - CDF User documentation created by Help-Desk members

SAM Deployment on FNAL CAFs

- First round of SAM deployment (SAM v6) at CDF for onsite computing is complete, and a success
- Second round – upgrade to newer version of SAM (v7) is in progress
- Created a test platform for load-testing SAM with CDF use patterns
  - used extensively by CDF DH group and CD SAM team
  - good communication on test results and decisions
- Adopted a methodical cycle of development-testing-deployment
- Deploying a new data-handling system on a running experiment is a challenge – we have demonstrated success and have a strategy to ensure continued success with SAM upgrades

New Sam-Farm for Data Reconstruction

- Sam-Farm concept and design overcomes certain limitations of previous Farm architecture
  - Successfully reviewed by CD in December 2004
  - New, more 'open' design using computing and data-handling tools already in use, such as SAM, CAF, dCache, Enstore ("just another CAF")
  - Analysis and reconstruction farms nearly identical now
  - Processing capacity can be increased/decreased easily
    - due to flexible boundary between different farms
  - Better error recovery => smoother operations
- Migrate from expert-only system to broader support base
  - Goal in 1 year – run by shift crew with experts consulting as needed

Sam-Farm Status

- Successfully deployed and commissioned during January-May 2005
- Now in production mode – all 2005 data processed by Sam-Farm
- Demonstrated
  - High-efficiency, high-throughput performance (see Aidan's talk)
  - Reassigning analysis farm CPU to reconstruction farm for short-term increase in reconstruction rate (motivated by need to catch up on 2005 data)
- Reconstruction farm is now integrated into a CDF-wide common computing platform

FronTier Status

- FronTier is a joint project with CD – a technology to distribute database access
- Becomes important as more computing (Monte Carlo production and user analysis) moves offsite – otherwise database access becomes a limitation.
- FronTier validation (obtaining identical analysis results with / without FronTier) was man-power limited, has recently been concluded.
- Freeze of CDF software release 'publishing' FronTier is imminent.

Fermi-Grid

- Fermi-Grid is a very important step in CDF's path to full GRID compatibility
  - Demonstrates interoperability with other Fermilab experiments on shared resources
- Fermi-Grid has good access to CDF data, making it an ideal platform for user physics analysis
  - User analysis is the dominant use case for data access
- Italian colleagues have demonstrated success with “condor glide-in” technique
  - Submitting CDF MC production and user analysis jobs to CPU not belonging to CDF, via one headnode
  - User interface identical to CAF
  - Fermi-Grid “glideCAF” is operational and in beta-testing

Increases in Offsite Computing

- Existing dCAF sites moving to shared pools (away from dedicated CDF resources)
  - eg. Italian dCAF no longer has any CDF-specific CPU
  - CDF owns a condor glide-in headnode
  - All jobs are running on LCG Tier 1 site
- New dCAFs starting with OSG / LCG shared pools
  - Short-term method: condor glide-in using one headnode per site
  - New sites coming up or under negotiation
    - Paris group (Lyon center)
    - Wisconsin GRID Laboratory
    - Chicago ATLAS Tier 2 site

Offline Operations

Data Processing

- Goal announced last year: move away from "built-in data reprocessing" mode (process all data for ~ 1 year, finalize calibrations for physics, reprocess all the data)
- Achieved "one-pass processing mode" in 2005
  - Preprocess calibration datasets only, turn around all calibrations in 4-6 weeks, use for official dataset production
  - Caught a few operator errors in the first two calibration cycles, increased validation scrutiny
- Continue to make incremental improvements to automate operations

MC Production

- Emphasize distributed operations, both CPU and personnel
  - Use offsite dCAFs (almost) exclusively
  - Job submission and disk buffer maintenance responsibility lies with physics group representatives
  - MC production group provides centralized software (eg. Tarball creation) support
- Best use case for pioneering GRID submission

Moving to the Future

Central Services

- Succeeded in creating a unified computing model for all applications (data reconstruction, MC, user analysis)
- Focus on the "Central Services" model
  - Plan to consolidate the infrastructure support further
- Emphasize and improve monitoring capability
  - Provide adequate coverage of all computing systems
  - Allow efficient triage and diagnosis of system problems
- Recruiting help from collaboration to work on monitoring
- Push to reduce operational load in the out-years

Moving to the GRID

- Data reconstruction uses fully self-contained tarball
- Plan to make MC production tarball also fully self-contained
- User analysis is a less well-defined problem
  - But is solvable in principle and we plan to move in this direction

Moving to the GRID

- Ongoing efforts (mainly by CDF Italian collaborators) to make the CAF interface operable with the LCG
  - Fruitful collaboration between CDF and LCG efforts in Italy
- Goal: user jobs submitted to "LCG-CAF" would be routed by LCG Resource Broker to the computing element
- Building on existing infrastructure, CDF is interested in similar path towards an OSG-CAF – needs more skilled people
- CDF and D0 joined forces (together with colleagues from UCSD, ANL, Boston, Chicago, Duke) to submit OSG preproposal
  - support SAM-Grid and Condor Glide-in within OSG framework

Next Step towards GRID

- Availability of more analysis CPU will help 1/fb challenge for winter and summer 2006
- Some of our offsite installations (eg. Italy, Wisconsin) are investing into disk as well as CPU
- With sufficient local disk cache we can target specific datasets to be served by remote dCAF
  - eg. considering a B physics analysis center at Italian dCAF, consistent with
    - Local group interest
    - Larger CPU per unit data needed for B analysis relative to high p_T physics
- Opportunity to exploit distributed data-handling capability of SAM

Summary

- Data processing and Monte Carlo generation schemes streamlined for fast turnaround
- Reconstruction and Simulation software mature and stable
- SAM and SAM-Farm are working well
- Increasing access to global computing resources to match physics needs